

DisCSPs with Privacy Recast as Planning Problems for Utility-based Agents

Julien Savaux,
Julien Vion,
Sylvain Piechowiak,
and René Mandiau
University of Valenciennes

Toshihiro Matsui
Nagoya Institute of Technology

Katsutoshi Hirayama
Kobe University

Makoto Yokoo
Kyushu University

Shakre Elmane, Marius Silaghi
Florida Institute of Technology

Abstract—Privacy has traditionally been a major motivation for decentralized problem solving. However, even though several metrics have been proposed to quantify it, none of them is easily integrated with common solvers. Constraint programming is a fundamental paradigm used to approach various families of problems. We introduce Utilitarian Distributed Constraint Satisfaction Problems (UDisCSP) where the utility of each state is estimated as the difference between the expected rewards for agreements on assignments for shared variables, and the expected cost of privacy loss.

Therefore, a traditional DisCSP with privacy requirements is viewed as a planning problem. The actions available to agents are: communication and local inference. Common decentralized solvers are evaluated here from the point of view of their interpretation as greedy planners. Further, we investigate some simple extensions where these solvers start taking into account the utility function. In these extensions we assume that the planning problem is further restricting the set of communication actions to only the communication primitives present in the corresponding solver protocols. The solvers obtained for the new type of problems propose the action (communication/inference) to be performed in each situation, defining thereby the policy.

I. INTRODUCTION

In Distributed Constraint Satisfaction Problems (DisCSP), agents have to find values to a set of shared variables while respecting given constraints (frequently assumed to have unspecified privacy implications). To find such assignments, agents exchange messages until a solution is found or until some agent detects that there is no solution to the problem. Thus, commonly agents reveal information during the solution search process, causing privacy to be a major concern in DisCSPs [22].

The artificial intelligence assumption is that utility-based agents are able to associate each state with a utility value [17]. As such each action is associated with the difference between initial and final utilities. If users are concerned about their privacy, then such a user can associate a utility value with the privacy of each piece of information in the definition of their local problem. Since the users are interested in solving the problem, they must be also able to quantify the utility each of them draws from finding the solution. Here we approach the problem by assuming that privacy has a utility that can be aggregated with the utility value of solving the problem. We evaluate how much privacy is lost by the agents during

the problem solving process, by the total utility of each information that was revealed. The availability of a value from the domain of a variable of the DisCSP in the presence of the constraints of an agent, is the kind of information that the agents want to keep private. For example, proposing an assignment with that value has a cost quantifying the desire of the agent to maintain its feasibility private. In traditional algorithms, agents participate in the search process until an agreement is found. We investigate the case where, being utility-driven, an agent may stop its participation if the utility of the privacy expected to be lost overcomes the reward for finding a solution of the problem. Simple extensions to basic algorithms are investigated to exploit the utilitarian model of privacy.

We then evaluate and compare synchronous and asynchronous algorithms according to how well they preserve privacy. To do so, we generate distributed meeting scheduling problems, as described in [16], [8]. In these problems, all agents own one variable, corresponding to the meeting to schedule, and the domain is the same for all variables. The constraints consist in a global constraint that requires all the variables to be equal, and also a unary constraint for each agent.

In the next section we discuss previous research concerning privacy for DisCSPs. Further we formally define the concepts involved in UDisCSPs. In Section IV we introduce some extensions to common DisCSP solvers that exploit the utilitarian model for privacy. After a discussion on theoretical implications, we present our experimental results in Section V. We conclude in Section VI.

II. BACKGROUND

A. Backtracking Algorithms

1) *Synchronous Backtracking*: The baseline algorithm for DisCSPs is the Synchronous Backtracking (SyncBT), as presented in [23]. SyncBT is a simple distribution of the standard backtracking algorithm. The agents start by determining a hierarchy between them. The higher priority agent then sends a satisfying assignment of its variable to the next agent with an *ok?* message. The recipient adds to it an instantiation of its own variable while respecting its constraints, and continues likewise. If an agent is unable to find an instantiation

compatible with the current partial assignment it has received, the agent sends a `nogood` message to the previous agent in the hierarchy. The process repeats until a complete solution is built, or until the whole search space is explored. The main efficiency concern is that, since the messages are being sent sequentially, it does not take advantage of possible parallelism.

2) *Asynchronous Backtracking*: Asynchronous Backtracking (ABT) [23], allows agents to run concurrently. Each agent finds an assignment of its variable and communicates it to the others agents, having constraints involving this variable. Agents then wait for incoming messages. They receive an `ok?` message containing an assignment from a related higher priority agent, at the beginning of the resolution and also each time such an agent changes its assignment to avoid constraint violation.

An agent eventually receives values proposed by the agents it is connected to by incoming links. These values form a context called `agent view`. When an agent receives an `ok?` message, it integrates the received assignment into its `agent view` and checks whether its own solution is consistent with it. If it is not the case, the agent's assignment is changed. The negation of a subset of an `agent view` preventing an agent from finding an assignment that does not violate any of its constraints, is called a `nogood`. If an agent infers a `nogood` from its constraints and its `agent view`, the assignment of the lowest priority agent involved in the `nogood` has to be changed. A `nogood` message communicates to that agent the `nogood`, which is treated by its recipient as a new constraint and can cause it to change its assignment and generate corresponding `ok?`, `addlink()` or `nogood` messages.

B. Privacy

In air traffic control [12], each airport has to allocate each takeoff and landing slots to the different flights. Even if airlines need combinations of slots to operate sequences of flights, slots are currently allocated individually. Such coordinated decisions are often impossible because of the need to keep constraints private [5]. Thus, privacy has been an important aspect for DisCSP solving algorithms. Recently, privacy preserving algorithms have also been developed for solving distributed constraint optimization problems [10] [19].

In existing works, there are two main approaches to enforce privacy. The first one uses cryptographic techniques. The main problem of these methods is that cryptographic protocols can be much slower, which often makes them impractical [11]. The second approach is based on using different search strategies to minimize privacy loss, as defined by certain privacy metrics. In this section we exemplify methods using these approaches.

1) *Sample Cryptographic Technique*: As an example of cryptographic technique, the approach described in [24], achieves a high level of privacy using encryption, giving more importance to privacy than to the efficiency of the resolution. It consists of using randomizable public key encryption scheme. In this algorithm, three servers (value selector, search controller and decryptor) receive encrypted information from

agents and cooperate to find an encrypted solution. Relevant parts of the solution are then sent to each agent. This method guarantees that no information is leaked to other agents. It also guarantees that, thanks to the renaming of values by permutation, servers cannot know the actual values they are dealing with. We now investigate methods that do not use cryptography.

2) *Distributed Private Constraint Satisfaction Problems*: A framework called Distributed Private Constraint Satisfaction Problems (DisPrivCSPs), is introduced in [7], [18], modeling the privacy loss for individual revelations. It also models the effect of the privacy loss by assuming that agents may abandon when the incremental privacy loss overcomes the expected gains from finding a solution. Each agent pays a cost if the feasibility of some of its tuple is determined by other agents. The reward for solving the problem is given as a constant. Those concepts were so far used for evaluating qualitatively existing algorithms, but were not integrated as heuristics in the search process. Privacy and the cost/utility usual optimization criteria of Distributed Constraint Optimization Problems are merged in [4] into a unique criterion.

3) *Valuations of Possible States*: The Valuations of Possible States (VPS) framework [15], [14], [9] measures privacy loss by the extent to which the possible states of other agents are reduced [21]. Privacy is interpreted as a valuation on the other agents' estimates about the possible states that one lives in. During the search process, agents propose their values in an order of decreasing preference. At the end of the search process, the difference between the presupposed order of preferences and the real one observed during search determines the privacy loss: the greater the difference, the more privacy has been lost.

4) *Partially Known Constraints*: The Partially Known Constraints (PKC) model [2], uses entropy, as defined in information theory, to quantify privacy and privacy loss. In this method, two variables x_1 and x_2 owned by two different agents may share a constraint. However, not all the forbidden couples (x_1, x_2) are known by both agents. Each agents only knows a subset of the constraints. During the search process, assignment privacy is leaked through `ok?` and `nogood` messages, like in standard algorithms. This problem is solved by not sending the value that is assigned to a variable in a `ok?` message, but the set of values compatible with this assignment. For `nogood` messages, rather than sending the actual assignments, an identifier is used to specify the state of the resolution and is used to check if some assignments are obsolete or not.

III. CONCEPTS

The Distributed Constraint Satisfaction Problem (DisCSP) is the formalism commonly used to model constraint problems distributed between several agents. It is represented by a quadruplet $\langle A, V, D, C \rangle$ where:

- A : a set of agents.
- V : a set of variables, each one being owned by a distinct agent.

- D : a set of domains, each of them defining available values for the corresponding variable.
- C : a set of constraints, each constraint being a relation imposed between two variables (i.e, $x_1 = x_2$).

An agent that reveals an assignment to another agent, incurs a cost. Once the information is revealed, we consider that it becomes public, meaning that revealing it to yet another agent will not degrade its privacy.

Example 1. Suppose a meeting scheduling problem between a professor and two students. They all consider to agree on a time to meet on a given day, to choose between 8am, 10am and 2pm. For simplicity, in the next sections, we will refer to these possible values by their identifier: 1, 2, and 3. The Professor A_1 is unavailable at 2pm, Student A_2 is unavailable at 10am, and Student A_3 is unavailable at 8am.

There can exist various reasons for privacy. For example, Student A_2 does not want to reveal the fact that he is busy at 10am (because he secretly took a second job at that time). The value Student A_2 associates with not revealing the 10am unavailability is the salary from the second job (\$2,000). The utility of finding an agreement for each student is the stipend for their studies (\$5,000). This is an example of privacy for absent values or constraint tuples.

Further Student A_3 had recently boasted to Student A_2 that at 8am he interviews for a job, and he would rather pay \$1,000 than to reveal that he is not. This is an example of privacy for feasible values of constraint tuples.

Thus, corresponding agents associate a cost of 1 to the revelation of their availability at 8am, equals to 2 for the one at 10am, and equals to 4 for the one at 2pm. The reward from finding a solution is 5.

a) *DisCSP*: The DisCSP framework models this problem with:

- $A = \{A_1, A_2, A_3\}$
- $V = \{x_1, x_2, x_3\}$
- $D = \{\{1, 2, 3\}, \{1, 2, 3\}, \{1, 2, 3\}\}$
- $C = \{x_1 = x_2 = x_3, x_1 \neq 3, x_2 \neq 2, x_3 \neq 1\}$

As it can be observed, DisCSPs cannot model the details regarding privacy considerations. Now we will show how existing extensions model the remaining details.

With DisPrivCSPs the additional parameters are P , to specify the privacy coefficient of each value, and R , to specify the rewards of each coefficient.

- $P = \{P_{A_1}, P_{A_2}, P_{A_3}\} = \{(1, 2, 4), (1, 2, 4), (1, 2, 4)\}$
- $R = \{R_{A_1}, R_{A_2}, R_{A_3}\} = \{(5), (5), (5)\}$

As we see, this framework successfully models all the information described in the initial problem and also measures the privacy loss for each agent. However, it was not yet investigated what is the impact of the interruptions when privacy loss exceeds the revenue threshold, or how agents could use these information to modify their behavior during the search process to preserve more privacy.

b) *VPS*: For this problem, with the VPS framework, the 3 participants have to suppose an order of preference between

all different possible values for each other agent. As agents initially do not know anything about others agents but the variable they share a constraint with, they have to suppose an equal distribution of all possible values for all other agents, meaning that they do not expect the feasibility of any value to be less secret, and so proposed first. In this direction one needs to extend VPS to be able to also model the kind of privacy introduced in this example.

c) *PKC*: With PKC, the individual unavailabilities are only known by the corresponding participant. Only the junction of information known by the two agents over a given constraint can reconstruct the whole problem.

- $A = \{A_1, A_2, A_3\}$
- $V = \{x_1, x_2, x_3\}$
- $D = \{\{1, 2, 3\}, \{1, 2, 3\}, \{1, 2, 3\}\}$
- $C = \{\{x_1 = x_2 = x_3, x_1 \neq 3\}, \{x_1 = x_2 = x_3, x_2 \neq 2\}, \{x_1 = x_2 = x_3, x_3 \neq 1\}\}$

Extensions of PKC can also be proposed to model our example by adding extra parameters for specifying the quantitative information about privacy, as shown below. Next we introduce a framework that can both specify the quantitative input details, and can help agents in their search process.

d) *UDisCSP*: While some previously described frameworks do model the details of our example, it has until now been an open question as to how they can be dynamically used by algorithms in the solution search process.

We propose to recast a DisCSP as a planning problem. It can be noticed that the rewards and costs in our problem bear similarities with the utilities and rewards commonly manipulated by planning algorithms [13]. As such, we propose to define a framework which, while potentially being equivalent in expressing power to existing DisCSP extensions, would nevertheless explicitly specify the elements of the corresponding family of planning problems.

We introduce the Utilitarian Distributed Constraint Satisfaction Problem (UDisCSP). Unlike previous DisCSP frameworks, besides results, we are also interested in the solution process. A policy is a function that associates each state of an agent with an action that it should perform [17].

We define an *agreement* as a set of assignments for all the variables with values from their domain, such that all the constraints are satisfied.

Definition 1. A UDisCSP is formally defined as a tuple $\langle A, V, D, C, U, R \rangle$ where:

- $A = \langle A_1, \dots, A_n \rangle$ is a vector of n agents
- $V = \langle x_1, \dots, x_n \rangle$ is a vector of n variables. Each agent A_i controls the variable x_i .
- $D = \langle D_1, \dots, D_n \rangle$ where D_i is the domain for the variable x_i , known only to A_i , and a subset of $\{1, \dots, d\}$.
- $C = \{c_1, \dots, c_m\}$ is a set of interagent constraints.
- $U = \{u_{1,1}, \dots, u_{n,d}\}$ is a matrix of costs where $u_{i,j}$ is the cost of agent A_i for revealing whether $j \in D_i$.
- $R = \langle r_1, \dots, r_n \rangle$ is a vector of rewards, where r_i is the reward agent A_i receives if an agreement is found.

The state of agent A_i includes the subset of D_i that it has revealed, as well as the achievement of an agreement. The problem is to define a set of communication actions and a policy for each agent such that their utility is maximized.

Note that the solution of a UDisCSP does not necessarily include an agreement. In principle the set of available actions for agents consist in any communication operator, as well as any local inference computation.

Example 2. The DisCSP in the Example 1 is extended to UDisCSP by specifying the additional parameters U, R :

$$U = \{u_{1,1} = 1, u_{1,2} = 2, u_{1,3} = 4, \\ u_{2,1} = 1, u_{2,2} = 2, u_{2,3} = 4, \\ u_{3,1} = 1, u_{3,2} = 2, u_{3,3} = 4\}. \\ R = \langle 5, 5, 5 \rangle.$$

The participants are utility-based agents [17] and try to reach the optimal state.

IV. ALGORITHMS

Now we discuss how the basic ABT and SyncBT algorithms are adjusted to UDisCSPs. The state of an agent includes the agent view. After each state change, each agent computes the estimated utility of the state reached by each possible action, and selects randomly one of the actions leading to the a state with the maximum expected utility.

In our algorithms, an information used by agents in their estimation of expected utilities is the risk of one of their assignments being rejected. This risk can be re-evaluated at any moment based on data recorded during previous runs on problems of similar parameters (e.g. problem density).

The learning can be online or offline. For offline learning one calculates the number of messages `ok?` and `nogood` sent during previous executions, called *count*. It also counts how many messages previously sent lead to the termination of the algorithm, in variable *terminationCount*. It calculates the risk for a solution to not lead to the termination of the algorithm, called *futilityRisk*. Alternatively one can update the *count*, *terminationCount* and *unsolvedRisk* dynamically whenever the corresponding events happen.

$$futilityRisk = 1 - \frac{terminationCount}{count} \quad (1)$$

When `ok?` messages are sent, the agent has the choice of which assignment to propose. When a `nogood` message is scheduled to be sent, agents also have choices of how to express them. Before each `ok?` or `nogood` message, the agents check which available action leads to the highest expected utility. If the highest expected utility is lower than the current one, the agent announces failure. The result is used to decide the assignment, `nogood`, or failure to perform.

We called these modified algorithms SyncBTU and ABTU, respectively. The algorithms SyncBTU and ABTU are obtained by performing the above mentioned modifications, in the pseudocodes of SyncBT [23], [25] and of ABT [22].

SyncBTU is obtained by restricting the set of communication actions to the standard communication acts of SyncBT, namely `ok?` and `nogood` messages. The procedures of a solver like SyncBT define a policy, since they uniquely identify a set of actions (inferences and communications) to be performed in each state. A state of an agent in SyncBT is defined by an agent-view and a current assignment of the local variable. The local inferences in SyncBTU are obtained from the ones of SyncBT by a simple extension exploiting the utility information available. The criteria in this research was not to guarantee an optimal policy but to use utility with a minimal change to the original behavior of SyncBT when reinterpreted as a policy. In SyncBTU, the state is extended to also contain a history of revelations of one's values defining an accumulated privacy loss, and a probability to reach an agreement with each action. Similar modifications are done to ABT to obtain ABTU: the restricted set of communication of ABTU is composed of `ok?`, `addlink` and `nogood`. The state and local inferences of ABTU are the same as SyncBTU, while also containing the set of `nogoods`.

For ABTU, there are three procedures of ABT that have to be modified: `checkagentview`, when `nogood`, and `backtrack`. The new procedure `checkagentview` is shown in Algorithm 1 and is obtained by inserting Lines 7 to 10. They test the privacy loss and only continue as usual if the expected loss is smaller than the expected reward.

For lack of space, we do not include here the modified versions of the other two procedures of ABTU, since they are obtained in the same way from the procedures of ABT in [22], procedure when `nogood`, 7th line, and procedure `backtrack`, 7th line. For SyncBTU, its procedures are obtained from the procedures of SyncBT in an identical way as for ABTU and ABT. Since [23] does not provide pseudocode for SyncBT, we refer the modifications to the pseudocode presented in [25], function `assignCPA`, before Line 7, and function `backtrack`, before Line 6.

Algorithm 1: procedure `checkAgentView` in ABTU

Input: $D, agentView, futilityRisk, reward$
Output:

```

1 when agentView and currentValue inconsistent do
2   if no value in  $D$  is consistent with agentView then
3      $\perp$  backtrack;
4 else
5   select  $d \in D$  where agentView and  $d$  are consistent;
6   currentValue =  $d$ ;
7   if  $calculateCost(futilityRisk, D, 1) \geq reward$ 
8     then
9        $\perp$  interruptSolving();
9   else
10     $\perp$  send(ok?, ( $x_i; d$ )) to outgoing links
```

To calculate the estimated utility of pursuing an agreement (revealing an alternative assignment), the agent considers all

different possible scenarios of the subsets of values that might have to be revealed in the future based on possible rejections received, together with their probability (see Algorithm 2). The algorithm assumes as parameters:

- the previously calculated *futilityRisk* (see Equation 1),
- the possible values D , and
- the probability of having to select from D .

The algorithm then recursively calculates the utility of the next possible states, and whether the revelation of the current value v leads to the termination of the algorithm, values stored in variables *costRound* and *costNonTerminal*. The algorithm returns the estimated cost of privacy loss for the future possible states currently, called *estimatedCost*.

Algorithm 2: CalculateCost

Input: *futilityRisk*, D , *probD*

Output: *estimatedCost*

```

1 if only one value is left in the domain then
2   return (marginalCost(value) * probD);
else
3    $v = \text{first}(D)$ ;
4    $\text{costRound} = \text{calculateCost}$ 
      ( $1 - \text{futilityRisk}$ ,  $\{v\}$ , probD);
5    $\text{costNonTerminal} = \text{calculateCost}$ 
      (futilityRisk,  $D \setminus \{v\}$ ,
      futilityRisk * probD);
6    $\text{estimatedCost} = \text{costRound} + \text{costNonTerminal}$ ;
7   return estimatedCost;

```

Example 3. Continuing with Example 1, at the beginning of the computation agent A_1 has to decide for a first action to perform. We suppose the *futilityRisk* learned from previous solvings is 0.5. To decide whether it should propose an available value or not, it calculates the corresponding *estimatedCost* by calling Algorithm 2 with parameters: the learned *futilityRisk* = 0.5, the set of possible messages ($D = \{1, 2, 3\}$) and *probD* = 1.

For each possible value, this algorithm recursively sums the cost for the two scenarios corresponding to whether the action leads immediately to termination, or not. Given privacy costs, the availability of three possible subsets of D may be revealed in this problem: $\{1\}$, $\{1, 2\}$, and $\{1, 2, 3\}$. The *estimatedCost* returned is the sum of the costs for all possible sets, weighted by the probability of their feasibility being revealed if an agreement is pursued. At the call, $\text{costRound} = u_{1,1} * 0.5$. In the next recursion for *costNonTerminal*, we get $\text{costRound} = (u_{1,1} + u_{1,2}) * 0.25$. In the last recursion, the algorithm returns $(u_{1,1} + u_{1,2} + u_{1,3}) * 0.25$. The *estimatedCost* obtained is $u_{1,1} * 0.5 + (u_{1,1} + u_{1,2}) * 0.25 + (u_{1,1} + u_{1,2} + u_{1,3}) * 0.25$. The expected utility ($\text{reward} + \text{estimatedCost} = 5 - 3 = 2$) of pursuing a solution being positive, the first value is proposed.

Next is an illustrative example of other ABTU operations.




Fig. 1. Interactions between agents during SyncBT

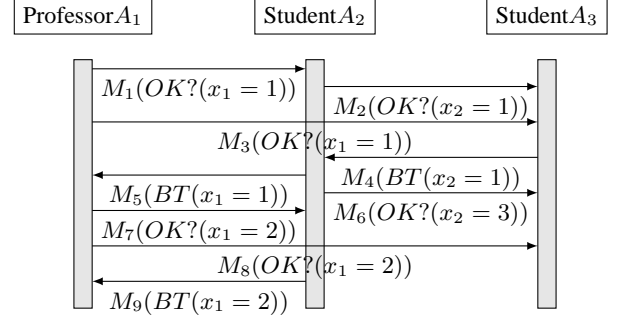


Fig. 2. Interactions between agents in ABT

Example 4. With the original SyncBT and ABT, possible obtained traces are depicted in Figure 1 and Figure 2, respectively. In Figure 2, we see that Student A_2 proposes $x_2 = 1$ in message M_2 and $x_2 = 3$ in message M_6 . In this case, the privacy loss for Student A_2 is $u_{2,1} + u_{2,3} = -1 - 4 = -5$.

However, with ABTU, we do not only use the actual utility of the next assignment to be revealed, but estimate privacy loss using Algorithm 2. After Student A_2 has already sent $x_2 = 1$ with M_2 , it considers sending $x_2 = 3$ with M_6 . This decision making process is depicted in Figure 3. If the next value, $2pm$, is accepted, Student A_2 will reach the final state while having revealed $x_2 = 1$ and $x_2 = 3$, for a total privacy cost of $u_{2,1} + u_{2,3} = -1 - 4 = -5$. If it is not, the unavailability of the last value $x_2 = 2$ will have to be revealed to continue the search process, leading to the revelation of all its assignments for a total cost of -7 . Since both these scenarios have a probability of 50% to occur, the *estimatedCost* equals $((-5 - 7)/2) = -6$. The utility ($\text{reward} + \text{estimatedCost}$) being equal to $5 - 6 = -1$, Student A_2 has no interest in revealing $x_2 = 3$ and interrupts the solving. Its final privacy loss is only $u_{2,1} = -2$. The utility of the final state reached by Student A_2 being -2 with ABTU, and -4 with ABT, ABTU preserves more privacy than ABT in this problem.

e) *Theoretical Discussion:* The introduced UDisCSP framework can assume without significant loss of generality that interagent constraints are public. This is due to the fact that any problem with private interagent constraints (e.g., PKC), is equivalent with its dual representation where each constraint becomes a variable [1]. However, note that privacy of domains mentioned in [3] is not modeled by privacy of




Fig. 5. Evaluation of privacy loss on instances with different parameters.

TABLE I
GENERAL COMPARISON FOR ALGORITHMS ALONG MULTIPLE METRICS

	SyncBT	ABT	SyncBTU	ABTU
PrivacyLoss	2,5	9.0	1,8	5,3
Messages	2,8	531,3	2,3	150,6
Solved	0,3	0,2	0,3	0,2
Interruption	0	0	0,3	0,7
cpuTime(ms)	257,7	1329,1	254,7	910,3

f) *Discussion on Experiments:* For each algorithm, we have measured in Table I the privacy loss, the number of messages exchanges, the number of problem solved, the number of solvings interrupted to preserved privacy and CPU time. The Figure 4 shows that synchronous algorithms are better than asynchronous at preserving privacy. Moreover, SyncBTU and ABTU are better than SyncBT and ABT at preserving privacy, respectively. The Figure 5 shows the averaged estimation of privacy loss per agent, according to the density and the distribution of the unary constraints. We observe that for a density higher than 0.4, privacy loss is lower for *tail-constrained DMS* than it is for *uniform DMS*, particularly with ABTU. We notice in Table I that interrupting the solvings to preserve privacy let agents not only reduce privacy loss by 39% but also reduce the calculation times by 27%, reduce the number of messages exchanged by 71% while still solving 98% of the problems solved by standard algorithms, the interruptions happening mostly when the problems have no solution.

VI. CONCLUSIONS

While many frameworks have been developed recently for coping with privacy in distributed problem solving, none of them is widely used, likely due to the difficulty in modeling common problems. In this article we propose a framework

called Utilitarian Distributed Constraint Satisfaction Problem (UDisCSP). It models the privacy loss for the revelation of an agent's constraints as a utility function. We present algorithms that let agents estimate how much privacy will be lost at the end of the solving process, using information from previous experience in solving problems. This estimation is used to modify the agent's behavior. We then show how adapted synchronous and asynchronous protocols (SyncBTU and ABTU) behave and compare them on different types of distributed meeting scheduling problems. The comparison shows that SyncBTU can maintain more privacy on random problems, as compared to both ABTU and original versions ABT and SyncBT. Some families of problems with particular properties regarding privacy were also identified.

In future work, we want to investigate how much privacy is leaked during the solving of different classes of problems. We also plan to improve the way agents learn from previous experience, by using not only the density of the corresponding problems, but also the tightness, the number of variables or the number of interagent constraints they are involved in.

REFERENCES

- [1] F. Bacchus and P. Van Beek. On the conversion between non-binary and binary constraint satisfaction problems. In *Proceedings of the 15th National Conference on Artificial Intelligence*, Madison, Wisconsin, July 1998.
- [2] Ismel Brito, Amnon Meisels, Pedro Meseguer, and Roie Zivan. Distributed constraint satisfaction with partially known constraints. *Constraints*, 14(2):199–234, 2009.
- [3] Ismel Brito and Pedro Meseguer. Distributed forward checking. In *CP*, 2003.
- [4] Prashant Doshi, Toshihiro Matsui, Marius Silaghi, Makoto Yokoo, and Mark Zanker. Distributed private constraint optimization. In *Web Intelligence and Intelligent Agent Technology*, 2008. *WI-IAT'08. IEEE/WIC/ACM International Conference on*, volume 2, pages 277–281. IEEE, 2008.
- [5] Boi Faltings, Thomas Léauté, and Adrian Petcu. Privacy guarantees through distributed constraint satisfaction. In *Web Intelligence and Intelligent Agent Technology*, 2008. *WI-IAT'08. IEEE/WIC/ACM International Conference on*, volume 2, pages 350–358. IEEE, 2008.
- [6] Ferdinando Fioretto, William Yeoh, and Enrico Pontelli. Multi-variable agents decomposition for dcops to exploit multi-level parallelism. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1823–1824. International Foundation for Autonomous Agents and Multiagent Systems, 2015.
- [7] Eugene C Freuder, Marius Minca, and Richard J Wallace. Privacy/efficiency tradeoffs in distributed meeting scheduling by constraint-based agents. In *Proc. IJCAI DCR*, pages 63–72, 2001.
- [8] A. Gershman, A. Grubshtein, A. Meisels, L. Rokach, and R. Zivan. Scheduling meetings by agents. In *Proc. 7th International Conference on Practice and Theory of Automated Timetabling (PATAT 2008)*. Montreal (August 2008), 2008.
- [9] Rachel Greenstadt, Jonathan P Pearce, and Milind Tambe. Analysis of privacy loss in distributed constraint optimization. In *AAAI*, volume 6, pages 647–653, 2006.
- [10] Tal Grinshpoun and Tamir Tassa. A privacy-preserving algorithm for distributed constraint optimization. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 909–916. International Foundation for Autonomous Agents and Multiagent Systems, 2014.
- [11] Martin Hirt, Ueli Maurer, and Bartosz Przydatek. Efficient secure multi-party computation. In *Advances in Cryptology-ASIACRYPT 2000*, pages 143–161. Springer, 2000.
- [12] International Air Transport Association IATA. *Worldwide scheduling guidelines*. International Air Transport Association, 2005.
- [13] Sven Koenig, David Furcy, and Colin Bauer. Heuristic search-based replanning. In *AIPS*, pages 294–301, 2002.

- [14] Rajiv T Maheswaran, Jonathan P Pearce, Emma Bowring, Pradeep Varakantham, and Milind Tambe. Privacy loss in distributed constraint reasoning: A quantitative framework for analysis and its applications. *Autonomous Agents and Multi-Agent Systems*, 13(1):27–60, 2006.
- [15] Rajiv T Maheswaran, Jonathan P Pearce, Pradeep Varakantham, Emma Bowring, and Milind Tambe. Valuations of possible states (vps): a quantitative framework for analysis of privacy loss among collaborative personal assistant agents. In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 1030–1037. ACM, 2005.
- [16] Rajiv T Maheswaran, Milind Tambe, Emma Bowring, Jonathan P Pearce, and Pradeep Varakantham. Taking dcop to the real world: Efficient complete solutions for distributed multi-event scheduling. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pages 310–317. IEEE Computer Society, 2004.
- [17] Stuart Russell and Peter Norvig. *Artificial intelligence: a modern approach*. 2010.
- [18] Marius-Calin Silaghi and Boi Faltings. A comparison of distributed constraint satisfaction techniques with respect to privacy. In *Proceedings of the Distributed Constraint Reasoning Workshop at AAMAS*. Citeseer, 2002.
- [19] Tamir Tassa, Roie Zivan, and Tal Grinshpoun. Max-sum goes private. In *Proceedings of the 24th International Conference on Artificial Intelligence*, pages 425–431. AAAI Press, 2015.
- [20] Richard J Wallace and Eugene C Freuder. Constraint-based reasoning and privacy/efficiency tradeoffs in multi-agent problem solving. *Artificial Intelligence*, 161(1):209–227, 2005.
- [21] R.J. Wallace and E.C. Freuder. Constraint-based multi-agent meeting scheduling: Effects of agent heterogeneity on performance and privacy loss. In *DCR*, pages 176–182, 2002.
- [22] Makoto Yokoo, Edmund H Durfee, Toru Ishida, and Kazuhiro Kuwabara. The distributed constraint satisfaction problem: Formalization and algorithms. *Knowledge and Data Engineering, IEEE Transactions on*, 10(5):673–685, 1998.
- [23] Makoto Yokoo, Toru Ishida, Edmund H Durfee, and Kazuhiro Kuwabara. Distributed constraint satisfaction for formalizing distributed problem solving. In *Distributed Computing Systems, 1992., Proceedings of the 12th International Conference on*, pages 614–621. IEEE, 1992.
- [24] Makoto Yokoo, Koutarou Suzuki, and Katsutoshi Hirayama. Secure distributed constraint satisfaction: Reaching agreement without revealing private information. In *Principles and Practice of Constraint Programming-CP 2002*, pages 387–401. Springer, 2002.
- [25] Roie Zivan and Amnon Meisels. Synchronous vs asynchronous search on discsps. In *Proceedings of the First European Workshop on Multi-Agent Systems (EUMA)*, 2003.